arXiv:0908.2493v1 [cs.CG] 18 Aug 2009

Minimum Feature Size Preserving Decompositions

Greg Aloupis* Erik D. Demaine† Martin L. Demaine‡
Vida Dujmović§ John Iacono¶

Abstract

The *minimum feature size* of a crossing-free straight line drawing is the minimum distance between a vertex and a non-incident edge. This quantity measures the resolution needed to display a figure or the tool size needed to mill the figure. The *spread* is the ratio of the diameter to the minimum feature size. While many algorithms (particularly in meshing) depend on the spread of the input, none explicitly consider finding a mesh whose spread is similar to the input. When a polygon is partitioned into smaller regions, such as triangles or quadrangles, the *degradation* is the ratio of original to final spread (the final spread is always greater).

Here we present an algorithm to quadrangulate a simple n-gon, while achieving constant degradation. Note that although all faces have a quadrangular shape, the number of edges bounding each face may be larger. This method uses $\Theta(n)$ Steiner points and produces $\Theta(n)$ quadrangles. In fact to obtain constant degradation, $\Omega(n)$ Steiner points are required by any algorithm.

We also show that, for some polygons, a constant factor cannot be achieved by any triangulation, even with an unbounded number of Steiner points. The specific lower bounds depend on whether Steiner vertices are used or not.

1 Introduction

1.1 Problem and Motivation

This paper contains an analysis of planar polygon decompositions, focusing on the objective that among all new point-edge distances produced, none should

*Université Libre de Bruxelles, Brussels, Belgium. aloupis.greg@gmail.com
†Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. edemaine@mit.edu
‡Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. mdemaine@mit.edu
§Carleton University, Ottawa, Ontario, Canada. vida@cs.mcgill.ca
¶Polytechnic Institute of New York University, Brooklyn, New York, USA. jiacono@poly.edu

be much smaller than those in the input. Specifically, the *minimum feature size* of a noncrossing planar straight line drawing (a special case of which is the noncrossing set of line segments forming a polygon) is the minimum distance between two non-touching edges. Our goal is to bound the change in minimum feature size that results from decomposing a polygon into triangles, quadrangles, and so on.

Feature size was introduced in the context of meshing [7], where it influences the necessary mesh complexity that guarantees certain levels of quality. See also [1, 5].

Another motivating example is the 95-year-old algorithm by Lowry [6, 3] for finding a common dissection of any two polygons of equal area. This algorithm starts by triangulating the polygon, then using a dissection from 1778 to convert each triangle into a rectangle with a common height ϵ equal to half the minimum height of all triangles. The algorithm uses $O(A/\epsilon)$ pieces where A is the area of the polygon, and this bound is clearly the best possible for each individual piece. The issue is that ϵ depends on the choice of triangulation. We show that a triangulation without Steiner points could be forced to have ϵ arbitrarily small, even if the input polygon has a constant number of vertices and a constant *spread*. The spread is the ratio of *diameter* over minimum feature size—roughly, the ratio between the largest and smallest distances. We remind the reader that the diameter is the largest distance between any two points in the union of the edges. The term "spread" comes from the analogous measure for point sets [2].

Essentially, this paper can be thought of as an implicit study of meshing with the objective of preserving feature size, or equivalently, spread (meshing does not affect diameter).

Lowry's algorithm is an example of a *pseudopolynomial bound*. Quite common in problems with integer inputs, pseudopolynomial bounds are polynomial in the input size, n, and in the sum of the input integers, often denoted N. Subset Sum, for example, has a pseudopolynomial-time algorithm, but has no polynomial-time algorithm unless P = NP, being weakly NP-complete [4].

Pseudopolynomial bounds are also common in computational geometry, where problems have real inputs such as point coordinates. One approach is to assume that the real inputs are rationals, which includes most representations of real numbers on a digital computer, scale them all to become integers, and let N be the sum of the resulting integers. However, some geometries, especially the outputs of geometric algorithms, are impossible to represent with rational coordinates. For this reason, computational geometers generally prefer the *real RAM* model of computation, where each input can be any real number, and algorithms have the ability to do only basic arithmetic, radical, and sign computations on these real numbers.

The usual approach for a pseudopolynomial bound on a real RAM is to measure the effective resolution of the input using basic geometric quantities. A pseudopolynomial bound might be polynomial in the input size n and in the spread r. Note that, as with problems with (distinct) integers, $n \leq N$, here we have a (slightly weaker) bound of $n = O(r^2)$.

While many algorithms attain pseudopolynomial bounds, e.g., on running time, and thus depend on spread, few algorithms consider the spread of their output, as we have seen in our example. This omission becomes important when trying to chain algorithms together. Thus one can see the importance of our objective of obtaining a bounded increase in feature size when producing a decomposition.

1.2 Preliminaries

Definition 1. *For any planar straight line drawing G, the* minimum feature size, $\mathrm{mfs}(G)$, *is the minimum distance between a vertex and a non-incident edge. Let* $\mathrm{diam}(G)$ *denote the diameter of G. The spread,* $\mathrm{spread}(G)$, *is* $\frac{\mathrm{mfs}(G)}{\mathrm{diam}(G)}$.

We are interested in subdividing an n-sided polygon P into a planar straight-line drawing G such that the minimum feature size of G is as close as possible to that of P. We call the ratio $\frac{\mathrm{mfs}(P)}{\mathrm{mfs}(G)}$ the *degradation* of the decomposition of P into G. In this paper we bound the degradation in terms of n, for various common types of decompositions.

We will be looking at three types of decompositions, known as *triangulations*, where the interior of P is partitioned into triangles. See Figure 1. The most common decomposition of a polygon is the *classic triangulation*, where non-crossing chords are added between vertices of P, until the interior of P is partitioned into triangles. More general types of triangulations allow the addition of vertices (known as *Steiner* points) on or inside P. We will consider two such triangulations; a *proper* triangulation forbids the placement of Steiner points in the interior of any edge. In other words, no two edges incident to a common Steiner point are collinear. In the dual graph of the triangulation all vertices have degree at most 3. A *non-proper* triangulation simply partitions P into triangles, with no restrictions. Three types of *quadrangulations*, or in general k-rangulations, are defined analogously.

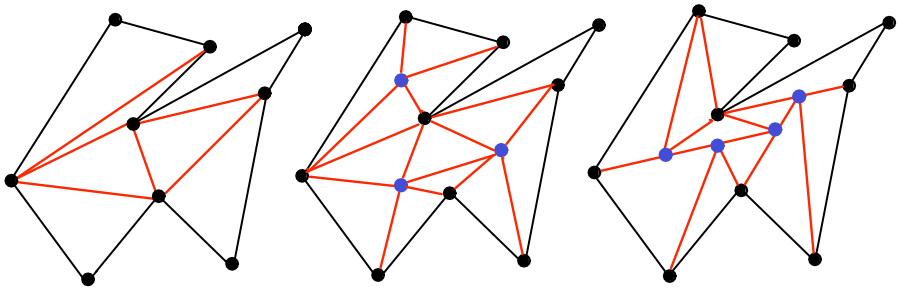

Figure 1: Types of triangulations: classic, proper, non-proper. Steiner points are blue.

	Triangulation	Quadrangulation
Classic	$\Omega(n)$	$\Omega(n)$
Proper	$\Omega\left(\frac{\log n}{\log\log n}\right), O(\log n)$	$\Omega\left(\frac{\log n}{\log\log n}\right), O(\log n)$
Non-proper	$O(\log n)$	$\Theta(1)$

Table 1: **Our Results.** The best minimum feature size degradation values are indicated, for various decompositions of an n-vertex polygon. The quadrangulation results also hold for any k-rangulation for $k \geq 4$. We conjecture that the three $O(\log n)$ values are tight.

1.3 Contributions in this paper

In section 2 we show that classic k-rangulations, for any k, cannot guarantee low degradation of minimum feature size. Specifically, we show all classic k-rangulations of a regular n-gon have a minimum feature size degradation of $\Omega(n)$. In Section 3 we show that there is a family of n-sided polygons for which all proper k-rangulations have a minimum feature size degradation of $\Omega\left(\frac{\log n}{\log\log n}\right)$. On the positive side, in section 4 we give a $\Theta(n)$-time algorithm that produces a $\Theta(1)$-degradation non-proper quadrangulation for any polygon with n sides.

2 Lower bound for classic triangulations

Lemma 1. *For every n, there is a polygon P_n such that all classic triangulations of P_n have degradation $\Omega(n)$.*

Proof. Consider a regular even n-gon P_n, with unit sides. The minimum feature size is 1. Thus the spread is $\sin\frac{\pi}{n}$.

By the *Two Ears Theorem*, any classic triangulation T of a polygon must contain two disjoint triangles, each containing two adjacent edges of the polygon. The minimum feature size of these triangles is $\sin\frac{\pi(n-2)}{2n}$. Thus any triangulation of P_n will have a degradation of $\Omega(n)$. □

The arguments of the above proof can be extended trivially to quadrangles or any decomposition of constant size.

3 Bounds on proper triangulations

In this section we prove that any constant-size decomposition (i.e. r-rangulation) of a certain class of n-gons must have a degradation of $\Omega(\frac{\log n}{\log\log n})$. We first bound the maximum degree that a vertex can have within a rectangle, in order to maintain a unit minimum feature size (section 3.1). The main result follows in section 3.2.

3.1 Maximum Degree in a Rectangle

Lemma 2. *Given vertex a and two incident edges $\overline{ab}$ and $\overline{ac}$, if minimum feature size $= 1$, then $\angle bac \cdot \min(|ab|, |ac|) \geq 1$.*

Proof. Refer to Figure 2. Assume w.l.o.g. that $1 \leq |ab| \leq |ac|$. The distance from b to $\overline{ac}$ is $|ab| \cdot \sin \angle bac$. This distance must be at least 1 in order to achieve unit minimum feature size. Since $|ab| \cdot \angle bac \geq |ab| \cdot \sin \angle bac$, our claim holds. □

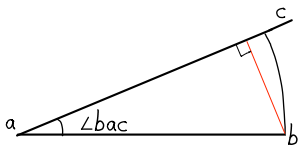

Figure 2: Proof of Lemma 2.

Definition 2. *For a point c in a rectangle R, let $d(c, \theta, R)$ be the distance from c to the boundary of R along a ray of angle θ with respect to the x-axis.*

The following lemma relies on elementary trigonometry and thus no proof is given.

Lemma 3. *Let c, p be two points in a $k \times n$ axis-aligned rectangle R. Let θ be the angle between $\overline{cp}$ and the x-axis. If $\theta \leq \tan^{-1} \frac{n}{k}$ then $d(c, \theta, R) \leq \frac{k}{\cos \theta}$. If $\theta \geq \tan^{-1} \frac{n}{k}$ then $d(c, \theta, R) \leq \frac{n}{\sin \theta}$.*

Theorem 1. *Within a $k \times n$ rectangle R, let a point c be connected to q points $\{p_1, \ldots, p_q\}$. The minimum feature size can equal 1 only if $q = O(k \log \frac{n}{k})$.*

Proof. Assume w.l.o.g. that R is axis-aligned and at least $q/4$ of the points p_i are above and to the right of c. Let this subset be denoted by P'. Let γ_i be the angle of $\overline{cp_i}$ relative to the x axis. Separate the points in P' into two groups P^ℓ, P^s depending on whether γ_i is larger or smaller (respectively) than $\tan^{-1} \frac{n}{k}$. Let P'' be the larger of the two groups. Let $l = |P''|$ and assume that the points in P are numbered such that P''=$\{p_1, p_2, \ldots p_l\}$. Furthermore for $1 \leq i < j \leq l$, if $P'' = P^s$, $\gamma_i \leq \gamma_j$ (otherwise if $P'' = P^\ell$, $\gamma_i \geq \gamma_j$).

We claim that $\int_{\gamma_i}^{\gamma_{i+1}} d(o, \theta, R) d\theta \geq 1$, where o is the lower-left corner of R. By Lemma 2:

$$1 \leq \angle p_i c p_{i+1} \cdot \min(|cp_i|, |cp_{i+1}|)$$

Since for all i the segment cp_i is inside R, $|cp_i| \leq d(c, \gamma_i, R)$, so

$$\leq \angle p_i c p_{i+1} \cdot \min(d(c, \gamma_i, R), d(c, \gamma_{i+1}, R))$$

For any given angle within the range considered here, the distance from o to the boundary of R is greater than the equivalent distance from c. Thus we have:

$$\leq \angle p_i c p_{i+1} \cdot \min(d(o, \gamma_i, R), d(o, \gamma_{i+1}, R))$$

Since the distances increase with the angles:

$$= \angle p_i c p_{i+1} \cdot d(o, \gamma_{i+1}, R)$$

Since $\angle p_i c p_{i+1} = \gamma_{i+1} - \gamma_i$:

$$= \int_{\gamma_i}^{\gamma_{i+1}} d(o, \gamma_{i+1}, R) d\theta$$

Again, because distances increase with angles:

$$\leq \int_{\gamma_i}^{\gamma_{i+1}} d(o, \theta, R) d\theta$$

Thus we establish our claim that $1 \leq \int_{\gamma_i}^{\gamma_{i+1}} d(o, \theta, R) d\theta$. Next we claim that the following holds:

$$8(q-1) \leq l = \sum_{i=1}^{l} 1 \leq \sum_{i=1}^{l-1} \int_{\gamma_i}^{\gamma_{i+1}} d(c, \theta, R) d\theta$$

$$\leq \int_{\gamma_1}^{\gamma_i} d(o, \theta, R) d\theta$$

The proof depends on whether $P'' = P^{\ell}$ (left) or $P'' = P^{s}$ (right). In the former case, angles range from 0 to $\tan^{-1} \frac{n}{k}$, whereas in the latter case they range from $\tan^{-1} \frac{n}{k}$ to $\frac{\pi}{2}$.

$$\leq \int_{0}^{\tan^{-1} \frac{n}{k}} d(o, \theta, R) d\theta \qquad \leq \int_{\tan^{-1} \frac{n}{k}}^{\frac{\pi}{2}} d(o, \theta, R) d\theta$$

By Lemma 3, the left and right inequalities respectively become:

$$= \int_{0}^{\tan^{-1} \frac{n}{k}} \frac{k}{\cos \theta} d\theta \qquad = \int_{\tan^{-1} \frac{n}{k}}^{\frac{\pi}{2}} \frac{n}{\sin \theta} d\theta$$

The right expression is simpler to bound so we handle it first. Since $n+k$ is a trivial upper bound on any distance in R,

$$\leq \int_{\tan^{-1} \frac{n}{k}}^{\frac{\pi}{2}} (n+k) d\theta$$

Having removed the dependency on θ, we manipulate the limits of the integral and use $\tan^{-1} x = \frac{\pi}{2} - \tan^{-1} x^{-1}$ to obtain:

$$= \int_0^{\tan^{-1}\frac{k}{n}} (n+k)d\theta = (n+k)\tan^{-1}\frac{k}{n}$$

Since $\tan^{-1} x = x - \frac{x^3}{3} + \frac{x^5}{5} - \frac{x^7}{7} - \frac{x^9}{9} \ldots$, and $n \geq k$, we get

$$= O\left((n+k)\frac{k}{n}\right) = O(k)$$

This concludes the right case. We return to the left case, $\int_0^{\tan^{-1}\frac{n}{k}} \frac{k}{\cos\theta} d\theta$. Since $\int \frac{1}{\cos\theta} d\theta = 2\tanh^{-1}\tan\frac{\theta}{2}$:

$$= 2k\tanh^{-1}\left(\tan\frac{\tan^{-1}\frac{n}{k}}{2}\right) - 2k\tanh^{-1}\left(\tan\frac{0}{2}\right)$$

Since $\tanh^{-1}\tan 0 = 0$:

$$= 2k\tanh^{-1}\tan\left(\frac{1}{2}\tan^{-1}\frac{n}{k}\right)$$

Since $\tanh^{-1} x = \frac{1}{2}(\log(1+x) - \log(1-x))$:

$$= 2k\frac{1}{2}\left(\log\left(1+\tan\left(\frac{1}{2}\tan^{-1}\frac{n}{k}\right)\right) - \log\left(1-\tan\left(\frac{1}{2}\tan^{-1}\frac{n}{k}\right)\right)\right)$$

Since $\tan\frac{1}{2}x \leq \tan x$ for $x \geq 0$:

$$\leq k\left(\log\left(1+\tan\tan^{-1}\frac{n}{k}\right) - \log\left(1-\tan\left(\frac{1}{2}\tan^{-1}\frac{n}{k}\right)\right)\right)$$

$$= k\left(\log\left(1+\frac{n}{k}\right) - \log\left(1-\tan\left(\frac{1}{2}\tan^{-1}\frac{n}{k}\right)\right)\right)$$

The value $\tan(\frac{1}{2}\tan^{-1}\frac{n}{k})$ ranges between 0 and 1. However it equals 1 only if $k = 0$, which is a degenerate case that we exclude. Therefore the second log term is well defined and always has a negative value. Since $n \geq k$, the negative value is bounded by a constant, so we can claim that:

$$\leq k\left(\log\left(1+\frac{n}{k}\right) + O(1)\right) = O\left(k\log\frac{n}{k}\right)$$

Combining the $O(k)$ and $O\left(k\log\frac{n}{k}\right)$ bounds for the left and right cases yields

$$q = O\left(k\log\frac{n}{k}\right)$$

□

3.2 Main Proof

Lemma 4. *Let T be a proper r-rangulation of polygon P. Let p be a vertex such that the leftward ray from p starts out within P and intersects some edge of T. Then p has a path of at most $r-2$ edges to a vertex at least 1 unit to its left.*

Proof. Refer to Figure 3. Let rq be the first segment crossed by the leftward ray from p in T. Assume w.l.o.g. that q is to the left of r. Note that since the minimum feature size of T is at least 1, the ray must travel at least a unit distance before hitting rq. Thus the horizontal distance from p to q is at least 1.

If T is a proper triangulation then p must be connected to q by an edge in T, as illustrated on the right in Figure 3. In general, if T is a proper r-rangulation, q and p must still be on the same face and are connected by a chain of at most $r-2$ edges. Note that this is not true in a non-proper k-rangulation.

□

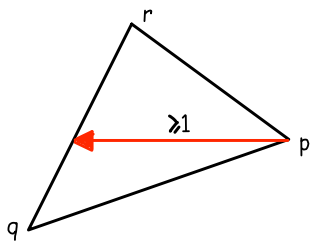

Figure 3: Illustration of Lemma 4.

Clearly, Lemma 4 holds regardless of the direction chosen for the ray. In general for a given direction, we say that the selected vertex q is the *parent* of p.

Given positive integers n and k, let $P(n,k)$ be the polygon illustrated in Figure 4. If its lower-left vertex is the origin, the five remaining vertices that define its shape have coordinates $(0,2k)$, $(k{+}nk^2,2k)$, $(k{+}nk^2,k)$, (k,k), $(k{+}nk^2,0)$. Thus the polygon is contained in a $(k{+}nk^2)\times 2k$ rectangle. An additional n vertices are placed k units apart from each other, starting from the top-left vertex on the top horizontal edge of $P(n,k)$.

The minimum feature size of $P(n,k)$ is k. We define the rectangle $R = [(k,k),(k,2k),(k{+}nk,2k),(k{+}nk,k)]$ as the *critical region*.

Theorem 2. *If $P(n,k)$ has a proper r-rangular decomposition with minimum feature size 1, then $k = \Omega\left(\frac{\log n}{\log\log n}\right)$.*

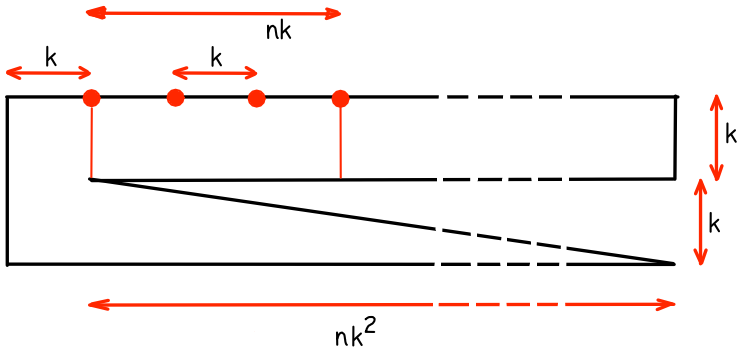

Figure 4: Lower bound example for proper r-rangulations: polygon $P(n,k)$.

Proof. First, observe that no Steiner points can be placed in the interior of the bottom edge of the critical region R, because they would be less than one unit away from the diagonal edge.

Since the decomposition is proper, each of the n vertices at the top of R must be incident to chords. These chords may end at the lower two vertices of R, or at Steiner points within R, or may lead outside of R via its left/right edges.

Notice that at most $k-1$ chords can cross either of the left/right edges of R, since the length of those edges is k.

Now consider the set of vertices in R, including any Steiner vertices added in a r-rangular decomposition. For each of these vertices, we choose the downward direction to construct parent relations, according to Lemma 4. The only vertices that are not assigned parents are the two located on the bottom edge of R. Thus the "sinks", i.e. vertices where the parent hierarchy ends, are the two bottom vertices and/or vertices outside R. These relations form a forest of trees, with a total of n leaves on the top edge of R. Note that the number of trees is at most $2k$ since we have established that this is the number of chords that can exit R. Thus one of these trees , T, must have at least $\frac{n}{2k}$ leaves.

According to Theorem 1, in a decomposition of a $nk \times k$ rectangle, the degree of any vertex must be $O\left(k \log \frac{nk}{k}\right) = O(k \log n)$ in order to preserve a unit minimum feature size. This implies that T has height $\Omega(\log_{k \log n}(n/2k))$. Furthermore, the same bound holds for the vertical distance from root to leaf, since each parent pointer represents a vertical separation of at least one unit, spread over at most $r-2$ vertices (where r is a constant).

However, by construction we know that this vertical distance is at most k. So we obtain $k = O(\log_{k \log n}(n/2k))$.

□

Theorem 2 tells us that by selecting $k = \frac{\log n}{\log \log n}$, the minimum feature size will drop from k to 1. Thus we obtain the following.

Corollary 1. *For any constant r, there exist n-gons for which all proper r-rangulations have a degradation of $\Omega\left(\frac{\log n}{\log \log n}\right)$.*

4 Non-proper quadrangulations can preserve minimum feature size

In this section we show how to construct a non-proper quadrangulation for any polygon, such that the minimum feature size degradation is $\Theta(1)$. We use $\Theta(n)$ Steiner points, and the construction can be computed in linear time.

Let $P=\{p_1, \ldots, p_n\}$ be an n-gon with minimum feature size 1. Consider the well-known *grassfire transformation*, $G_P(t)$, commonly used to visualize the formation of the medial axis of a polygon P. $G_P(t)$ is the result of shrinking P by lighting a fire along its boundary, and assuming that the fire will progress at unit speed and will last for t units of time. If P has unit minimum feature size, $G_P(t)$ will produce a connected region as long as $t < 0.5$.

We construct the curve $P_2 = G_P(\frac{2}{5})$. Once more we apply the grassfire transformation, this time to P_2, but we expand outward. In some sense we are computing $P_1 = G_{P_2}^{-1}(\frac{1}{5})$, where the inverse sign denotes outward expansion. It is not difficult to see that P_1 is contained in $G_P(\frac{1}{5})$. Thus we have constructed a "tube" of constant[1] width $1/5$, confined between $G_P(\frac{1}{5})$ and $G_P(\frac{2}{5})$.

Our tube has the property that any point inside it has distance at least $1/5$ to P. This is important because we intend to place Steiner vertices in the tube, specifically in the vicinity of each bend. In fact we will construct a polygon with boundary contained in the tube.

The tube consists of n bends, one per vertex of P. Each bend consists of a "greater" circular arc (spanning an angle of at most π), a "minor" co-circular arc (possibly a degenerate point, spanning the same angle at constant distance $1/5$ from the greater arc), and two segments of length $1/5$, joining the arc endpoints (at right angles). Let us call the two segments *doors*. Between each bend (from door to door), the tube has a rectangular shape of width $1/5$.

We will loosely construct our (Steiner) polygon, by providing a circular region for the location of each vertex. These circles will have a diameter of $1/20$, i.e. a quarter of the tube's width. Later on we will fix the precise location of each vertex in its region. Each circle will have its center on a specific curve which we call the tube's *track*. The track is constructed as follows. Within each bend, from door to door, the track is co-circular to the bend at a constant distance of $1/40$ from the greater arc. Between two bends, the track is a straight connecting segment.

Now we can fill the track with circles, evenly spacing them so that two consecutive circles are $1/20$ apart. The first and last circles added may be

[1] Here, constant width means that every point on the outer boundary has distance $1/5$ to some point on the inner boundary, and vice versa. In fact, only reflex vertices, where the tube bends, have multiple equidistant points (along circular arcs).

3/20 apart, but then we can re-position everything and the spacing will become slightly larger, without affecting our claims.

A "greater" arc has length at most $\frac{2}{5}\pi$, so the track within one bend has length at most $\frac{15}{40}\pi$. Thus the number of circles intersecting each bend is constant (roughly 25, at most. In fact this can probably be reduced to around 3-4). We delete all circles not intersecting a bend. Thus the total number of circles (and Steiner vertices) will be $O(n)$.

Now we superimpose our entire structure on a square grid of resolution 1/40. This means that every circle will contain at least one grid point. For each circle we choose such a point as a Steiner vertex. Next we join the Steiner vertices in the order of the circles appearing on the track. This creates a polygon $P' = \{a_1, a_2, \ldots, a_{(O_n)}\}$.

The size of the circles is small enough that any polygon connecting them will not intersect P_1. An equivalent statement is to say that the convex hull of two consecutive circles avoids P_1. This is trivially true for the rectangular components, and for any bend where the smaller arc is on P_2. On the other hand if the bend's greater arc is on P_2, the only edge of the hull that could intersect P_1 is the segment s tangent to the circles, closest to P_1. Consider the circular arc c that is 1/4 of the way from P_2 to P_1. The two circles fit tightly in the sub-tube between c and P_2. We place two segments, tangent to the two circles in question, and orthogonal to c. Let x_1 and x_2 be the intersection points of the segments with c. Every point on the segment between x_1 and x_2 is closer to P_1 than s is. Thus it suffices to establish that this segment does not intersect P_1. The arc corresponding to this segment is the angle spanned by three circles. As established earlier, this angle is a constant (roughly $\frac{3}{25}\pi$). Thus the closest point of the segment to the lesser arc of the bend is also a constant. Thus the separation between P and P' is at least 1/5.

Furthermore, it can be verified that no two Steiner vertices within a bend will have a separation smaller than that between consecutive vertices. What remains is to prove that no two bends are close to each other (in other words, we must prove that the Steiner vertices of P' are sufficiently separated). Every bend is associated with a vertex of P, or, to put it differently, with two consecutive edges of P. We know that every Steiner vertex of a bend has distance at most 2/5 from both associated edges of P. When we compare Steiner vertices from two different bends, we can identify one associated edge from each bend, such that the chosen edges are entirely disjoint (i.e. they don't even have a common endpoint). Since those two edges must have a separation of at least 1, we can conclude that the two groups of Steiner vertices have separation at least 1/5.

At this point, we can conclude that the union of P and P' has a minimum feature size of $\Theta(1)$. We can also add chords from each p_i to all the Steiner vertices in its associated bend. This creates several triangles that are roughly isosceles and have a common apex. Since the angle of each such triangle at the common apex is a constant, the separation between the newly introduced chords and their opposite vertices is also constant. As a result, the zone Z between P and P' becomes quadrangulated, with constant feature size.

Let T be the trapezoidal decomposition of P' resulting from extending hori-

zontal rays from each of its vertices. The construction of T places an additional $O(n)$ Steiner vertices on P' (i.e. on the inner boundary of Z). The minimum feature size of T is $\Theta(1)$ since all its vertices are on our grid of resolution $1/40$. Horizontal lines are $1/40$ apart, which means all new vertices are at least $1/40$ apart from other vertices, old or new. Since the new vertices are on P', they are at a safe distance from the boundary of P as well. Thus we conclude the following.

Theorem 3. *The planar straight line drawing consisting of the union of Z and T is a non-proper quadrangulation with a constant minimum feature size.*

Acknowledgments

The research presented in this paper was initiated at the 24$^{\text{th}}$ Annual Winter Workshop on Computational Geometry at the Bellairs Research Institute of McGill University. The authors would like to thank Godfried Toussaint as well as the non-author participants of the workshop for their helpful comments and for creating an environment conducive to creative thought: Zachary Abel, Brad Ballinger, Nadia Benbernou, Jit Bose, Jean Cardinal, Sebastien Collette, Mirela Damian, Marty Demaine, Robin Flatland, Ferran Hurtado, Scott Kominers, Stefan Langerman, Robbie Schweller, David Wood, and Stefanie Wuhrer.

References

[1] Tamal K. Dey. Delaunay mesh generation of three dimensional domains. Technical Report OSU-CISRC-09/07-TR64, Ohio State University, October 2007.

[2] Jeff Erickson. Nice point sets can have nasty delaunay triangulations. *Discrete & Computational Geometry*, 30(1), July 2003.

[3] Greg N. Frederickson. *Dissections: Plane and Fancy.* Cambridge University Press, November 1997.

[4] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., 1979.

[5] Benoit Hudson, Gary Miller, and Todd Phillips. Sparse voronoi refinement. In *Proceedings of the 15th International Meshing Roundtable*, pages 339–356, Birmingham, Alabama, September 2006. Springer-Verlag.

[6] Mr. Lowry. Solution to question 269, [proposed] by Mr. W. Wallace. In T. Leybourn, editor, *Mathematical Repository*, volume 3, part 1, pages 44–46. W. Glendinning, London, 1814.

[7] Jim Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 83–92, Austin, Texas, 1993.